Department of Physics
Computational Inference
Assignment 1 (due 5 Sept by close of play)

1. Consider the joint pdf $\pi(x, y, \theta)$ over x, y, and θ. Using Bayes' rule, product factoring, etc, establish the following identities between the pdfs over various conditional and marginal distributions:

   (a) $\pi(\theta|y) = \dfrac{\pi(x,\theta|y)}{\pi(x|\theta,y)}$(This is useful because the RHS may be evaluated at *any* x for which denominator is non-zero.)

   (b) $\pi(x,\theta|y) = \pi(x|\theta,y)\,\pi(\theta|y)$, i.e., we can condition the usual product factoring.

   (c) $\pi(\theta|y) = \dfrac{\pi(x,y|\theta)\,\pi(\theta)}{\pi(x|\theta,y)\,\pi(y)}$ (Again, RHS may be evaluated at any x for which denominator is non-zero. Useful when conditioning on θ is a GMRF.)

2. MatLab code to carry out inference by MCMC for the diffusion coefficient can be found on the module's web page https://coursesupport.physics.otago.ac.nz/wiki/pmwiki.php/Module412

   (a) Run that code (or python equivalent) and produce a plot of the posterior pdf over D conditioned on the (fake) measured data. Give a best value, and measure of uncertainty.

   (b) Plot the integrated autocorrelation time versus acceptance rate for this MCMC, by varying the proposal window (for fixed fake data).

   (c) Modify the formulation and code to allow the final time `Tmax` to be uncertain, uniformly distributed in the range $[1.9, 2.1]$. In particular, modify the code to perform joint inference over D and the (nuisance parameter) `Tmax`, and plot the posterior pdf over D conditioned on the (fake) measured data. Give a best value, and measure of uncertainty. How does the uncertainty compare to the result when `Tmax` was treated as certain?

3. Assume the $x_1, x_2, \ldots$ are a set of samples drawn from some distribution $\pi(x)$ (perhaps by an MCMC) and $y = g(x)$ is some function of x, with corresponding pdf $\pi(y)$ (given by the change of variable formula).

   (a) Show that $g(x_1), g(x_2), \cdots$ are a set of samples distributed as $\pi(y)$.

   (b) Use this property to plot distributions over the parameters x and $\sqrt{x}$ in the 'no such thing as best estimate' example from Lecture 1. In particular, set $x_* = 1$ and generate synthetic data

   $$d_i \sim N\left(\frac{\sqrt{x_*}}{i}, \sigma^2\right) \qquad i = 1, 2, \ldots N$$

   for some suitably large N and suitable σ^2. Then write an MCMC to draw samples from $\pi(x|\{d_i\})$, and plot a histogram of this distribution. Plot a histogram of $\pi(\sqrt{x}|\{d_i\})$. Report means, variances, etc. Repeat this for each of the prior distributions $\pi(x) \propto 1$ (i.e. perform inference using the normalized likelihood) and $\pi(x) \propto 1/x$ (the Jeffrey's prior for scale parameter).